7090 DATA PROCESSING SYSTEM

GENERAL INFORMATION

THE IBM 7090 Data Processing System, newest addition to IBM's family of data processing systems, includes the latest electronic component developments resulting from many years of accumulated experience in the data processing field. Main features of the system are:

- Increased speed in both internal processing and tape character access.
- Increased input-output flexibility with as many as eight data channels.
- Automatic priority processing, using data channel trap.
- New transistor circuits for greater speed and reliability.
- Program compatibility with the IBM 709 Data Processing System.

Incorporated in the 7090 are many of the outstanding capabilities of existing machines and a number of new features. This combination makes the 7090 a fast and efficient data processing system.

Data processing machines have become an integral part of many phases of science, business, and industry. Our rapidly expanding scientific investigations need faster methods for carrying out increasingly complex calculations. In addition, a vast amount of data is constantly being used in such areas as aircraft manufacture, government agencies, and retail establishments of all kinds. To meet these demands, the IBM 7090 Data Processing System has been developed.

A comparison of some of the features of the 7090 and 709 systems shows that:

- The 7090's internal processing speed is at least five times faster than that of the 709.
- The 7090 uses a maximum of eight input-output data channels. The 709 has no more than six.
- High-speed magnetic tape units may be used on the 7090.
- High-speed and low-speed tape units may be intermixed on the 7090.
- Both tape units operate in a high or low density mode.
- The 7090 costs less to install and operate than the 709.
- A separate operator's console is used with the 7090.

The use of transistors in the 7090 instead of vacuum tubes reduces both the total power and air conditioning requirements of the 7090 system by as much as 70 percent.

The modular unit design used in the 7090 results in a compact package. These modular units may be placed side by side, thus saving as much as 50 percent of the space required for the 709 system.

This bulletin assumes that the reader is familiar with the 709 system and its operation.

EXECUTION SPEEDS

Chart 1 shows average execution speeds for fixed or floating point arithmetic operations (including time needed to take a word from or put a word into core storage).

FIXED POINT OPERATIONS

Operation	709 Time and Op/Sec		7090 Time and Op/Sec	
Add or Subtract	24 usec*	41,660	4.8 usec	208,330
Logical Operations	24 usec	41,660	4.8 usec	208,330
Multiply	190 usec	5,260	27.8 usec	35,970
Divide	240 usec	4,160	33.6 usec	29,760

FLOATING POINT OPERATIONS

Operation	709 Time and Op/Sec		7090 Time and Op/Sec	
Add or Subtract	84 usec	11,900	14.6 usec	68,490
Multiply	170 usec	5,880	26.5 usec	37,730
Divide	216 usec	4,620	31.2 usec	32,050

Operations-per-second numbers are approximate.

*usec = microseconds

Chart 1. IBM 709 and IBM 7090 Average Execution Speeds

INFORMATION FLOW

It is characteristic of most applications that certain components of the data processing system are in operation almost constantly while others are called upon only occasionally. For example, the processing of a tape file to update the records affected by current transactions may keep two tape units and the data channels operating continuously, while the computing facilities and other components are used only a small portion of the time. In this case, it would be economical to use the remaining computing time to perform some other application not requiring frequent use of the data channels. On the other hand, another application might require extensive computation and yet make little reference to input-output devices. In this case, it would be desirable to use available facilities by performing simultaneously another application requiring little computing.

In the past, the problem of determining, through the stored program, when a component is available for use was so complex as to make simultaneous operation of two independent programs extremely difficult. Some computers overcome this problem by suspending all computing while an input-output device is being used. This is shown in Figure 1, where the central processing unit waits until the input device has accumulated a data word. This data word is then stored in core storage. Only then can the central processing unit begin to process it.

The problem of asynchronous balance is resolved in the 7090 system through a combination of two new units that make it possible to perform simultaneous reading, writing, and computing.

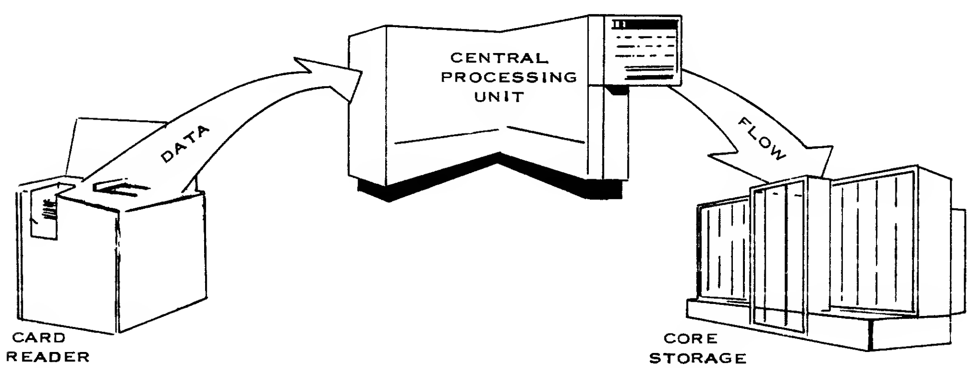

Figure 1. Data Transmission, 704 System

Multiplexor

One of these new devices is called the IBM 7606 Multiplexor. This unit accomplishes all of the data switching in the 7090 system. All components in the system must feed their data through the multiplexor; the only data path to core storage is from the multiplexor; likewise, any data coming from core storage must go first to the multiplexor.

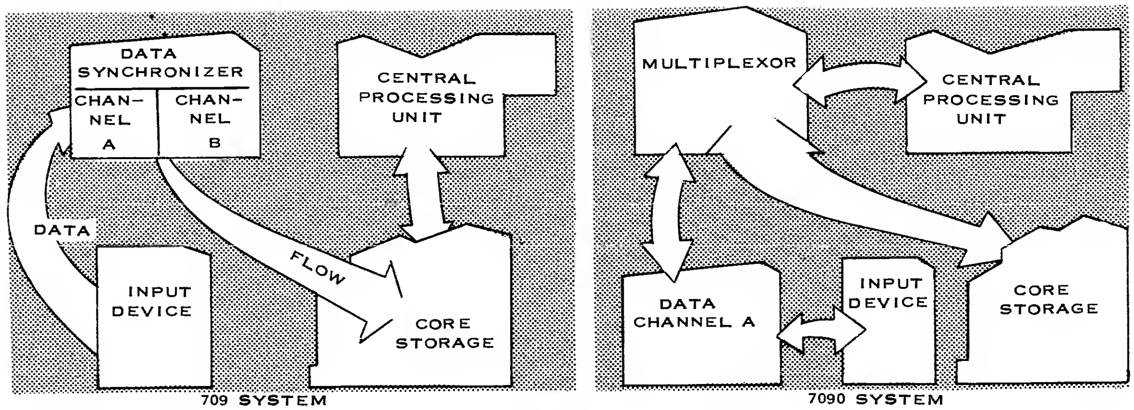

Figure 2. Data Flow in the 709 and 7090 Systems

A comparison of data flow in the 709 and 7090 systems is shown in Figure 2. In the 709 system, data flow is from an input-output device to a data synchronizer channel and then to core storage. The data synchronizer control section tells which data channel is to be used. In the 7090 system, however, the data flow path is from an input-output device to a data channel, to the multiplexor, and then to core storage. The data word is handled in the normal fashion of 36 bits in parallel, coming from or going to core storage.

Data Channel

The IBM 7607 Data Channel is another new unit used with the 7090 system. As many as eight data channels are accommodated by the multiplexor. Data channels replace the data synchronizers and tape control units used with the 709 system.

The 709 central processing unit performs a more remote role as an input-output controller than did the comparable unit in the 704 system. In the 7090 system the central processing unit's function is even more remote. It simply initiates and monitors input-output operations, but is not involved with the detail of routing data words or recognizing the ends of records or blocks of data. These functions are handled by the data channels and the multiplexor as shown in Figure 3.

Each data channel is basically a completely separate and independent input-output channel. The data channel provides for the transmission of data between the input-output units and core storage--a transmission that is independent of computing. A data channel may be thought of as a small computer with the responsibility for controlling the quantity and destination of all data transmitted between core storage and the input-output unit. It also performs limited counting and testing operations exclusively concerned with the transmission of data.

There are no restrictions on the type of input-output operation that can be performed by a data channel. With this type of computer system, the stored program starts an input operation by defining which input device is to transmit and which core storage area is to receive. The stored program may then proceed to execute instructions (compute), not related to the input transmission. When a full word has been assembled in the data channel, this word is put into core storage.

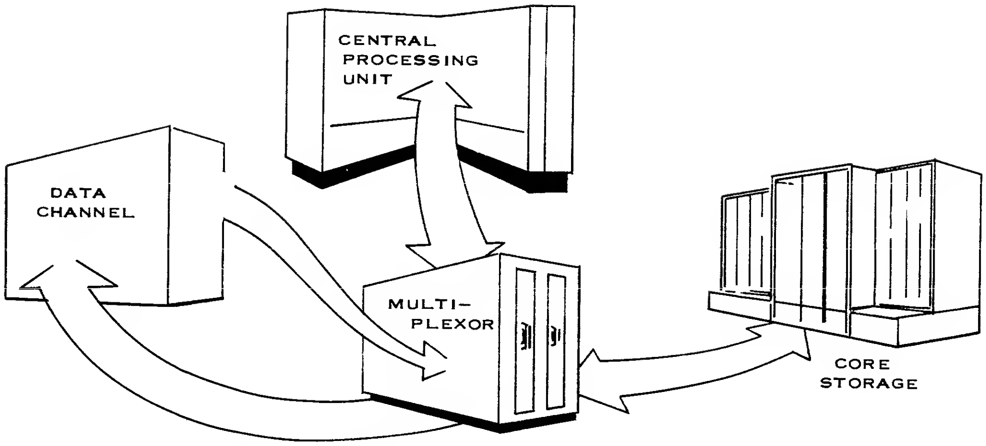

Figure 3. Data Transmission, 7090 System

This procedure is reversed if an output operation is being performed.

Thus, one or more data channels can be occupied with reading from magnetic tape while others are being used in writing on different tapes. Similarly, cards could be read or punched and results could be printed. All operations would be simultaneous with the processing that is occurring in the central processing unit.

Each data channel has its own separate console with all necessary keys and lights to assist the operator in entering or displaying data and commands. Also included is a unit select switch for changing the input-output unit being used.

A schematic of two data channels operating asynchronously with processing in the central processing unit is shown in Figure 4.

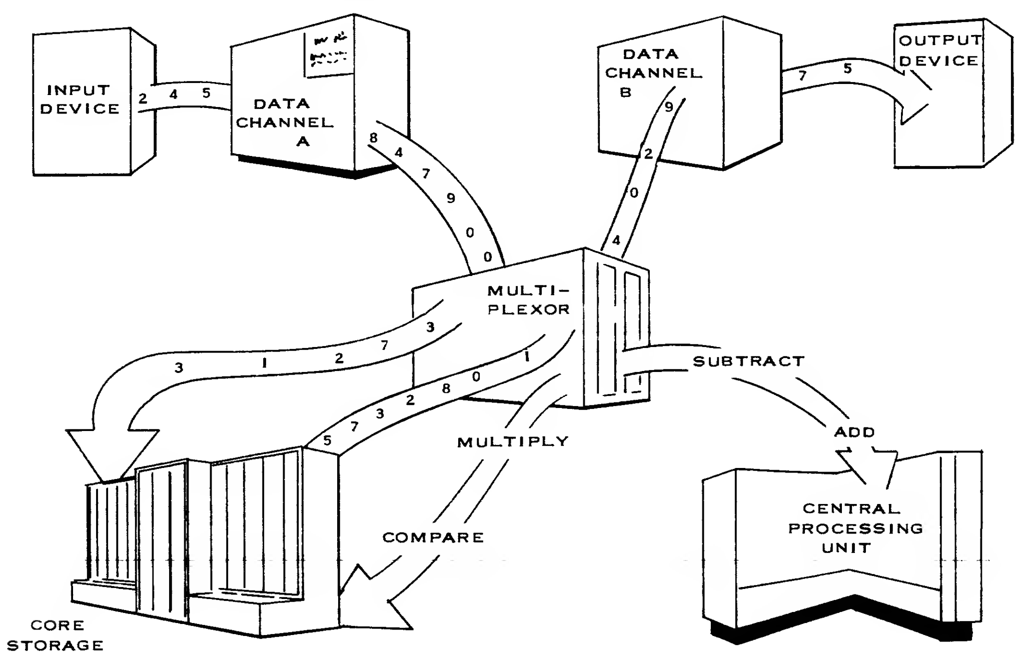

Figure 4. Data Channel and Central Processing Unit Data Flow

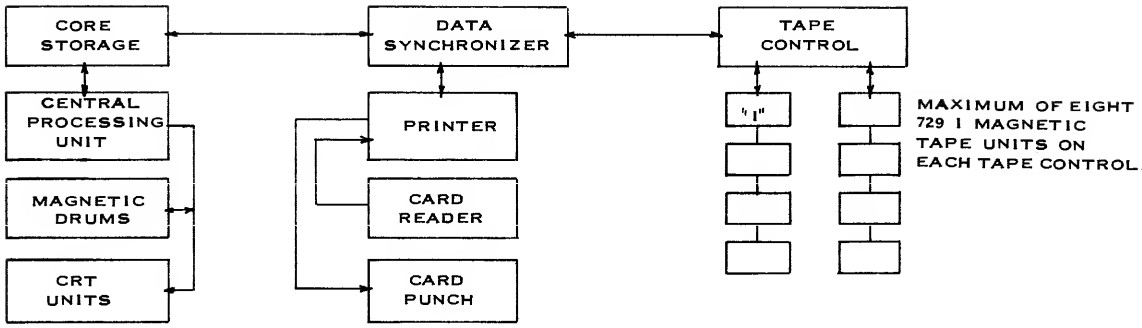

Figure 5. Information Flow, 709 System

A comparison of data flow in the 709 and 7090 systems is shown in Figures 5 and 6. With data being read by a 729 I tape unit (numbered 1 in Figure 5), the data would go first to the tape control, then to data channel A of the data synchronizer, and finally to the core storage unit being used with the system.

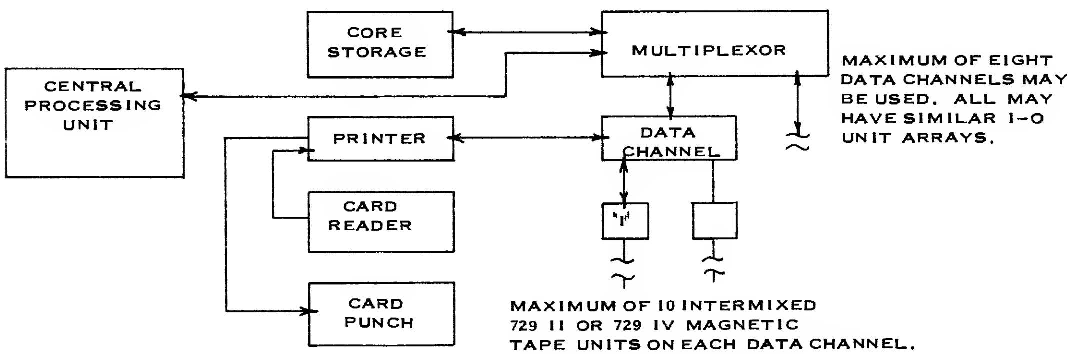

Figure 6. Information Flow, 7090 System

In Figure 6, the data flow is based on the assumption that data from core storage are being written on a 729 IV Magnetic Tape Unit (numbered 1). The data flow from the IBM 7302 Core Storage to the 7606 Multiplexor, to the 7607 Data Channel, and are then written on the magnetic tape unit.

This data channel system provides the features of the input-output control listed below. These items may be regarded as retention of and improvements on previous IBM methods of input-output control, together with the new ability of intermixing magnetic tape unit types. The items may be summarized briefly as the ability to:

1. Use the tape format of other input-output devices. This means that magnetic tapes from other IBM data processing systems may be used with the 7090 system, thus making it compatible with these systems.
2. Transmit input data to scattered core storage locations without additional programming (as in the 709 system).
3. Ignore unwanted data in a record, allowing the unwanted data to be skipped.
4. Use minimum time in transmitting data between input-output units and core storage. Only one storage cycle (2.4 usec) is used for each word processed.

5. Synchronize the stored program with input-output operations when desired. Synchronization is thus optional to the programmmer.
6. Allow the data channel to signal the central processing unit and save the current instruction's location whenever an end-of-file, a tape check, or a lack of a load channel instruction occurs.

(To these six refinements, add a seventh and new ability to:)

7. Operate eight input-output devices and compute simultaneously. From one to ten 729 II or 729 IV tape units may be attached to each data channel. The tape units may be intermixed in any fashion without addressing restrictions. Each data channel may also have a card reader, card punch, and a printer attached, in addition to the ten tape units.

Using all data channels, the maximum 7090 system includes: 80 intermixed 729 II and 729 IV tape units, eight card readers, eight card punches, and eight printers.

DATA CHANNEL TRAP

This feature controls an input-output program using the trapping concept. It provides the 7090 system with an extremely efficient and flexible means for data handling. The data channel is allowed to signal the central processing unit when an end-of-file or tape check occurs, or when certain data channel commands are executed and the accompanying load channel instruction is missing in the stored program.

When a trap occurs, the data channel is usually disconnected, the contents of the instruction counter are stored, and the next instruction is taken from a fixed core storage location (depending on the data channel being used).

Data channels may also be prevented from causing a trap until the program is able to handle it. In this event, the trapping signal is saved until the program allows it to become effective.

Data channel commands may employ indirect addressing. In this case, the address portion of the word in the location referred to by the instruction is used as the address in executing the command. Execution of an indirectly addressed command requires an additional storage cycle.

The entire operation of the data channel trap is described in the IBM 709-7090 Data Processing System Reference Manual, Form A22-6503.

MAGNETIC TAPE UNITS

The 7090 system uses both the 729 II and 729 IV tape units. The complete conversion to transistors in these units assures greater reliability. As with other 729 tape units, a two-gap recording mechanism is used, offering increased checking while writing. For magnetic tape advantages, see Figure 7.

Chart 2 points out the compatibility features of the different magnetic tape units and data processing systems.

1. Compact storage.

2. Fast reading and writing speeds.

3. Variable record size--from one to several thousand alphamerical characters per record.

4. Permanent files--reels may be used many times for many different jobs.

5. Lower cost--capacity of one reel is equal to several thousand IBM cards.

6. Compatible--prepared tapes may be used on many different IBM computer systems.

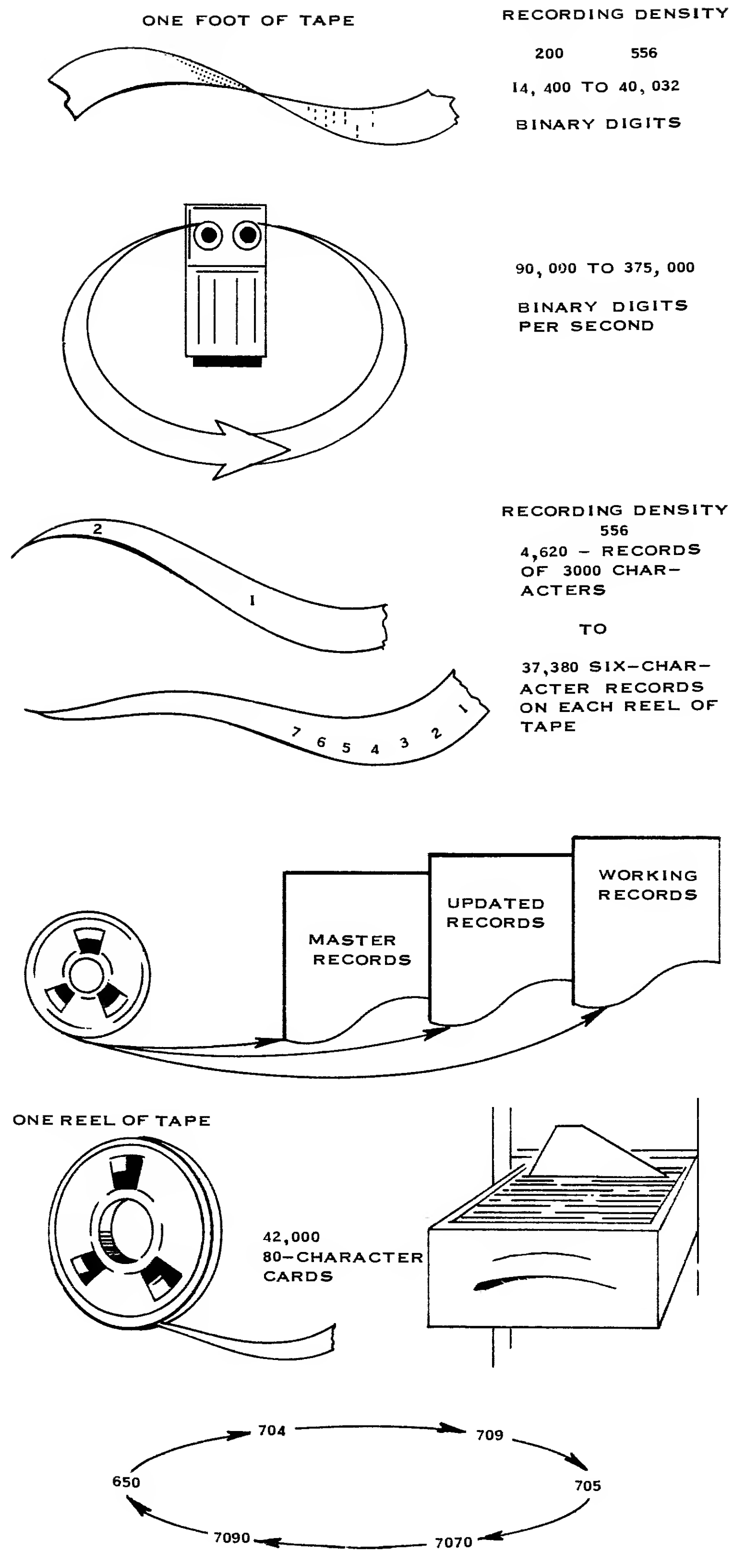

Figure 7. Magnetic Tape Advantages

		727	729 I	729 II	729 III	729 IV
Tape Speed (Inches per Second)		75	75	75	112.5	112.5
Recording Density (per Inch)		200	200	200 or 556	556	200 or 556
Character Rate per Second		15000	15000	15000 or 41667	62500	22500 or 62500
Record Gap Size in Inches		3/4	3/4	3/4	3/4	3/4
Record Gap Time in Milliseconds		10.8	10.8	10.8	7.3	7.3
Maximum Number of Tape Units per Control Unit:	IBM 753	10				
	IBM 754	10				
	IBM 755		8			
	IBM 760	2				
	IBM 767 (1)		5	and/or	5	
	IBM 777	8				
	IBM 7607 (2)			10	or	10
Where used:	Aux. Card-to-Tape (3)	x	x			
	Aux. Tape-to-Card	x	x			
	Aux. Tape-to-Printer					
	IBM 717	x	x			
	IBM 720-730	x	x			
	IBM 650	x				
	IBM 704	x				
	IBM 705 I or II	x				
	IBM 705 III		x		x	
	IBM 709		x			
	IBM 7070			x		x
	IBM 7090			x		x

(1) A maximum of ten 729 I tape units may be used if no 729 III tape units are used; otherwise, a maximum of five of each type must be used.
(2) Ten intermixed tape units may be used on each IBM 7607.
(3) Dual level sensing is not active.

Chart 2. Compatibility Features of IBM Magnetic Tape Units

Both the 729 II and 729 IV tape units read and write at one of two character densities; either 200 or 556 characters per inch. This dual-density feature is under stored program control. The combination of character density and tape speed provide character rates of 15,000 or 41,667 characters per second on the 729 II and 22,500 or 62,500 characters per second on the 729 IV tape unit.

A change density key is included with the other control keys on the 729 II and 729 IV tape units. Two new lights, high-density and low-density, on each tape unit indicate which character density is being used. Depression of the change density key simply changes the recording density.

Chart 3 compares the number of data records contained on tape reels written on low-speed and high-speed tape units.

Assuming a record length of 300 characters per record and a total of 63,200 records, a saving of two tape reels is achieved by using high-speed tape units (Figure 8).

	RECORDING DENSITY	
RECORD LENGTH	200	556
(No. of Char.)	(No. of Rec'ds/Reel)	(No. of Rec'ds/Reel)
6	36,460	37,375
30	31,600	35,373
60	27,085	33,069
300	12,640	22,046
1,800	2,916	7,127
3,000	1,805	4,624

NOTE: 2370 feet of usable tape with 3/4 inch record gaps is assumed. All figures are approximate.

Chart 3. Comparison of Capacities of Low-Speed and High-Speed Tape Unit Reels

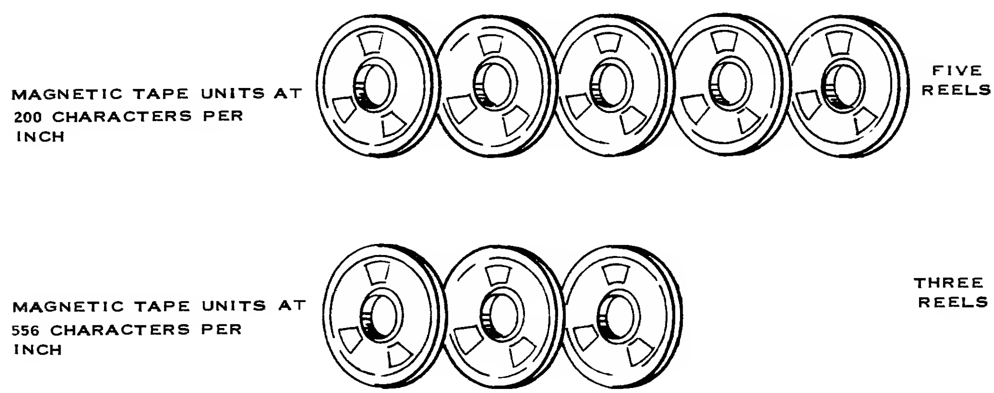

Figure 8. Comparison of High- and Low-Speed Tape Units

EXTERNAL SIGNAL

A standard feature of the computer system is its ability to accept a signal from an external source. This signal will cause the computer to execute a trapping operation when it is received. The instruction being executed is completed and the location of the next instruction in sequence is placed in the address part of core location 0003. The computer then takes its next instruction from location 0004.

DIRECT DATA FEATURE

As an optional feature, the computer may be equipped with the Direct Data feature. This allows the transmission of data between the computer and an external data device. With this feature, an external signal initiates the trapping operation and the stored program may then take whatever action is required to introduce data into the system or supply data from the system to the Direct Data device.

Thus the main office of a company using the computer could receive data from branch offices by direct wire. This direct data path, as is shown in Figure 9, could be an IBM 65-66 Transceiver or any one of many types of transmitting devices.

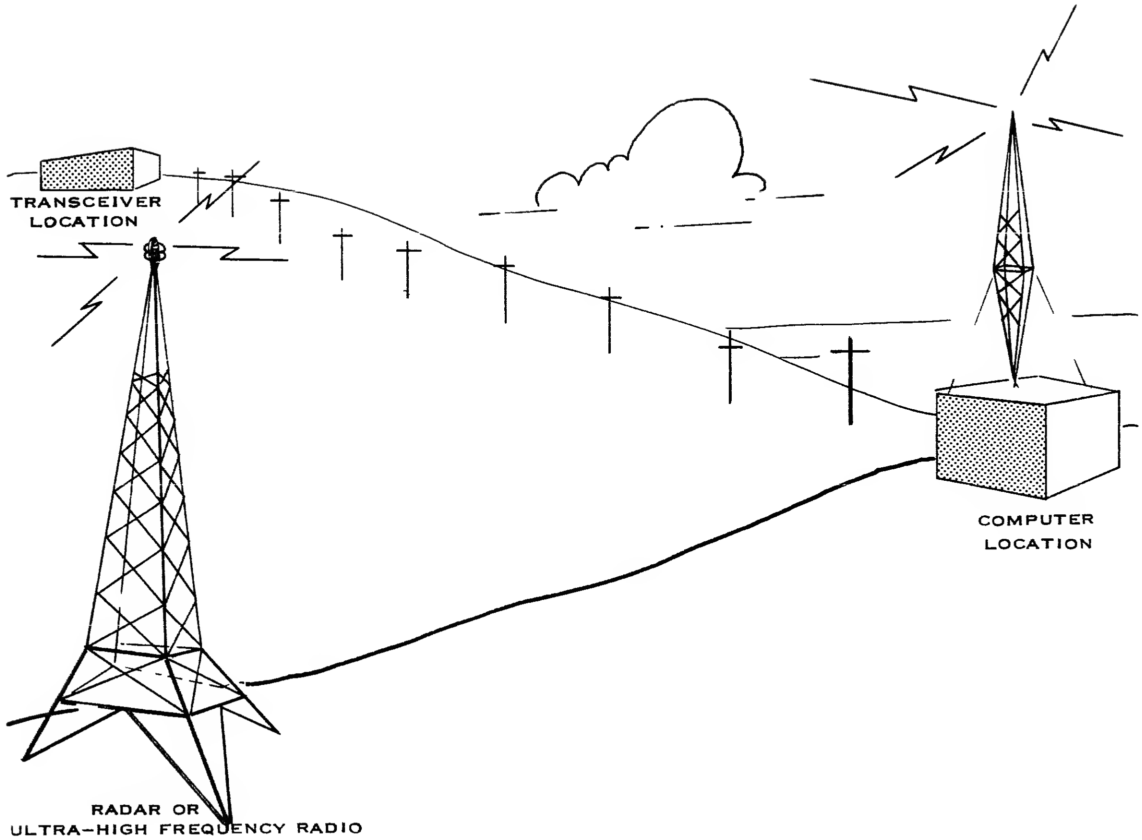

Figure 9. Schematic, External Signal and Direct Data Features

CONSOLE

The IBM 7151 Console Control is a separate unit providing centralized control of the 7090 system. It contains indicators, switches, keys and register displays for the operator's use. Channel indicators for data channel operation are provided and the register displays have been grouped for convenience. Data in any storage location can be displayed and/or changed manually by use of the entry keys and switches.

A maintenance feature, marginal checking device, is also a part of the console. This device is used to vary voltages and frequencies during programmed diagnostic testing, to detect potential difficulties before components actually fail.

INSTRUCTIONS AND PROGRAMS

More than 200 different operations may be accomplished with the 7090 system. All arithmetic and logical execution information is contained in the IBM 709-7090 Data Processing System Reference Manual. The execution of all instructions is the same; only the time required for them is different.

The 7090 system does not incorporate a magnetic drum or any cathode ray tube equipment. To insure that 704 and 709 programs will run on the 7090, the following principles should be observed:

1. Magnetic drum select instructions, copy, copy and add logical word, and the locate drum address instructions are executed as no-operation instructions and the input-output check indicator is turned on unless the computer has been previously set in the compatibility mode of operation.
2. Twice the core storage size of a given 704 program is required by the 7090 to process the program.

SHARE Programming System

The SHARE 709 system is the main programming system to be used with the 7090. It enables the programmers to write, check out, and alter their programs quickly and easily. The SHARE 709 system (SOS) can be thought of as composed of four distinct parts:

1. The SHARE compiler assembler translator (SCAT).
2. The program testing and correcting system.
3. The input-output system.
4. The MockDonald control system.

The SCAT portion of the SOS consists of two subparts, the compiler and the modify-and-load program.

The compiler performs about the same functions for the 709 and 7090 systems that the SHARE assembly program (SAP) performs for the 704 system. With a few minor exceptions, a SHARE assembly language program is acceptable as input for the compiler and results in a program listing and an absolute binary program card deck. In addition to the results produced by SHARE assembly program, the compiler can produce a "squoze" card deck. This deck contains the symbolic source program in encoded binary form. This form may be converted to machine language and loaded by the modify-and-load program almost as rapidly as an ordinary binary load program loads an absolute binary card deck. Two main reasons for the intermediate squoze card deck phase are:

1. Modifications to the program can be made in the original SCAT language and then added to the squoze deck for loading by the modify-and-load program.
2. Enough of the original symbolic information can be retained during the program execution to permit the checking and correcting program to give back printed output in the original symbolic language. These two features make it unnecessary, in most cases, to "think" or to "patch" a program in machine language. Thus, most or all of the cross-referencing between symbolic and binary codes may be avoided.

Another powerful tool in the compiler is "macro-operation." The compiler is built to recognize a large fixed set of macro-operations. It also accepts and temporarily retains definitions of macro-operations given by the programmer. In either case, it generates and inserts into the program the sequence of machine words specified by any one of these macro-operations in a macro-instruction. Among the system (fixed) macro-operations are all of the pseudo-operations making up the checking and correcting program, the input-output, and the MockDonald control system languages.

A significant feature of the SCAT system is that the loading process is also an assembly process. The squoze deck is not in a form that can be inserted into the computer "as is." It is the result of what corresponds to the first pass of the 704 SHARE assembly program system. With the SHARE assembly program, a second pass is needed to decode the squoze deck into absolute binary. In the SCAT system, the second pass is a function of the modify-and-load program. This program provides for the same modification of the original code that could be obtained by changing the original symbolic deck and reassembling. Thus, in SCAT, program modifications are given to the loading program and, for these changes, the modify-and-load program performs both the first and second passes. In this sense, the loading program is a full assembly program. When program modifications are made with the loader, the loader produces on request a new squoze deck or an absolute binary deck, as well as the listing of the modified program. In addition to the above operations, the modify-and-load program performs checking operations for both programmer and machine errors.

The input-output system permits writing of I-O programs designed for a particular customer's application. Transmission macro-instructions are used and are executed by a routine called the "dispatcher." These macro-instructions provide input-output that is simultaneous with computing. In addition, computing is interlocked with data transmission so that the computer will not attempt to use or modify data not yet complete. Transmission orders are channel stacked when required and subsequently the dispatching of these orders on data channels is automatic when the channel is free. The programmer may interrogate the "dispatcher" for the status of any transmission at any point in the program. The checking of input-output indicators is automatically accomplished by the "dispatcher."

The MockDonald control system has been designed to enable the automatic transition from one problem to the next, to maintain a machine program log, to aid in the parallel operation of tapes and CPU processing. In general, it performs many of the operations that would normally be handled by a professional machine operator. Chart 4 is a feature comparison of the 704, 709 and 7090 systems.

	704	709	7090
Core Storage Size	4096, 8192, or 32,768 words		32,768 words
Transistorized	no	no	yes
Internal Speed (Basic Cycle)	12 usec	12 usec	2.4 usec
Simultaneous Read-Write-Compute	no	yes	yes
Tape Skip Ability	no	yes	yes
Automatic Input-Output Priority	no	yes	yes
Data Channel Trapping	no	yes	yes
Number of Instructions (Approx.)	88	208	220
Maximum Number of:			
Tape Units	10	48	80
I-O Data Channels	none	2 through 6	1 through 8
Card Readers	1	3	8
Card Punches	1	3	8
Printers	1	3	8
Magnetic Drums	2	2	none
CRT Recorder and Display	1	1	none

Chart 4. Feature Comparison of the IBM 704, 709, and 7090 Systems

Refer to the IBM 709-7090 Data Processing System Reference Manual for the operation and instructional facts about the 7090 system. Except for the differences mentioned in this bulletin, the two systems are similar.

Form G22-6505-2 does not obsolete G22-6505-1 but contains new information on the following items:

1. Faster arithmetic execution times (chart 1)
2. Data channel console
3. Change density switch on tape units

IBM

International Business Machines Corporation
590 Madison Avenue, New York 22, New York

Printed in U.S.A. G22-6505-2 9/59